



(12) UK Patent (19) GB (11) 2 405 495 (13) B

(45) Date of publication: 20.09.2006

(54) Title of the invention: Data storage system

(51) INT CL: G06F 17/30 (2006.01)

(21) Application No:	0319393.5	(72) Inventor(s): Peter Bryan Malcolm
(22) Date of Filing:	18.08.2003	(73) Proprietor(s): Orchestra Limited (Incorporated in the United Kingdom) 190 The Strand, LONDON, WC2R 1JN, United Kingdom
(43) Date A Publication:	02.03.2005	
(52) UK CL (Edition X): G4A AMC		
(56) Documents Cited:		(74) Agent and/or Address for Service: Reddie & Grose 16 Theobalds Road, LONDON, WC1X 8PL, United Kingdom
EP 1168176 A2	EP 0351109 A3	
EP 0312865 A3	WO 2003/027891 A1	
WO 2002/101535 A2	WO 1994/018634 A1	
JP 2003006354 A	US 6378054 B1	
US 4310883 A	US 20030079084 A1	
US 20030050729 A1	US 20020147734 A1	
(58) Field of Search: As for published application 2405495 A viz: UK CL (Edition W) G4A INT CL ⁷ G06F, G11C Other: Online; EPODOC, JAPIO, WPI, INSPEC, TXTE. updated as appropriate		

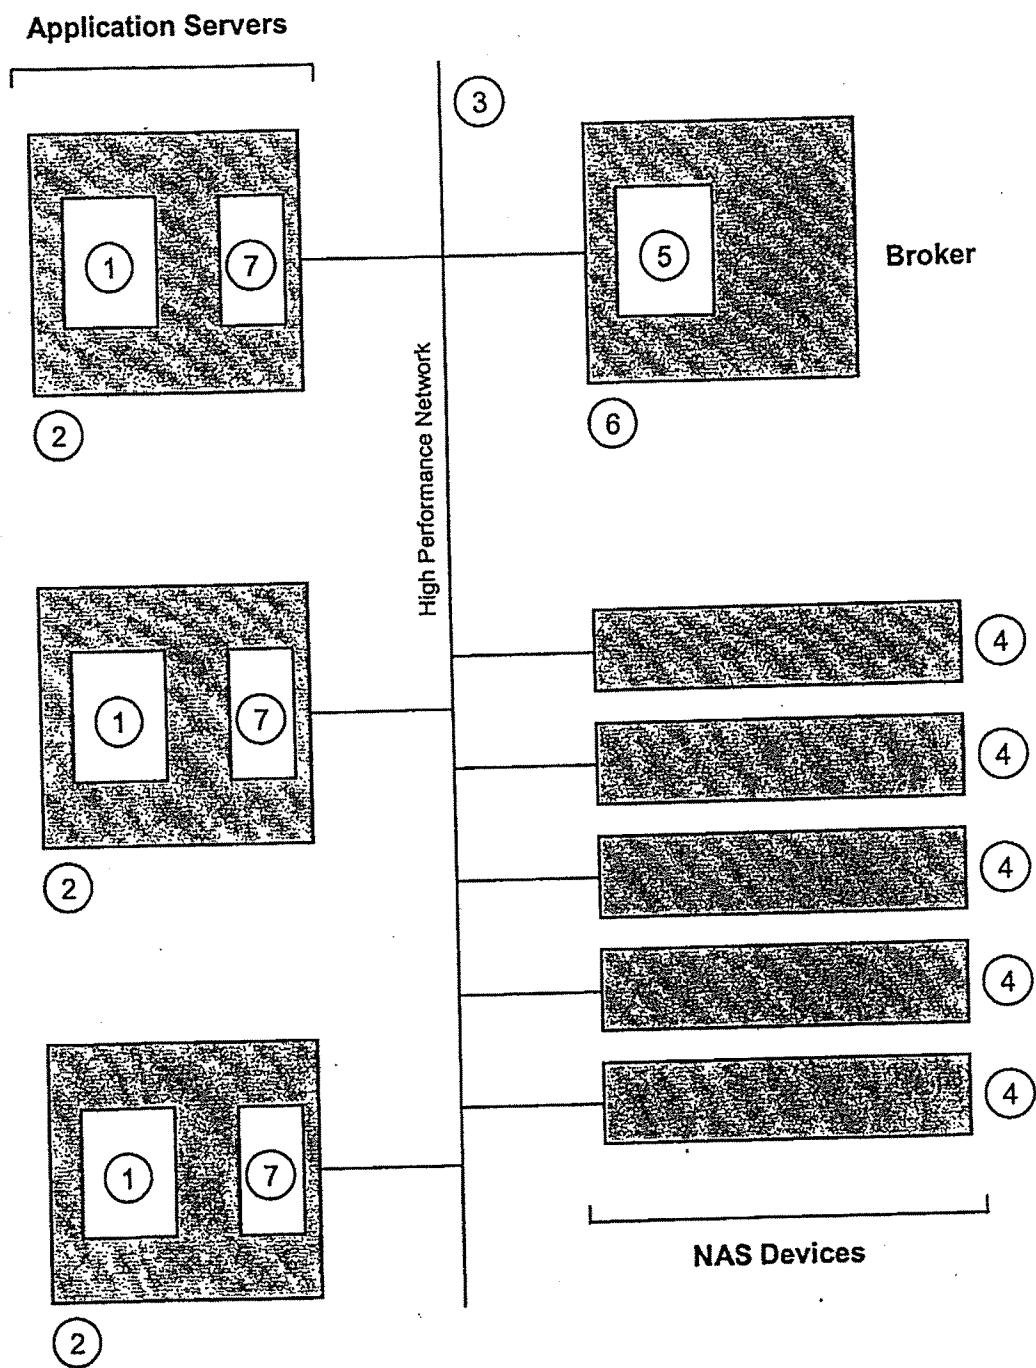
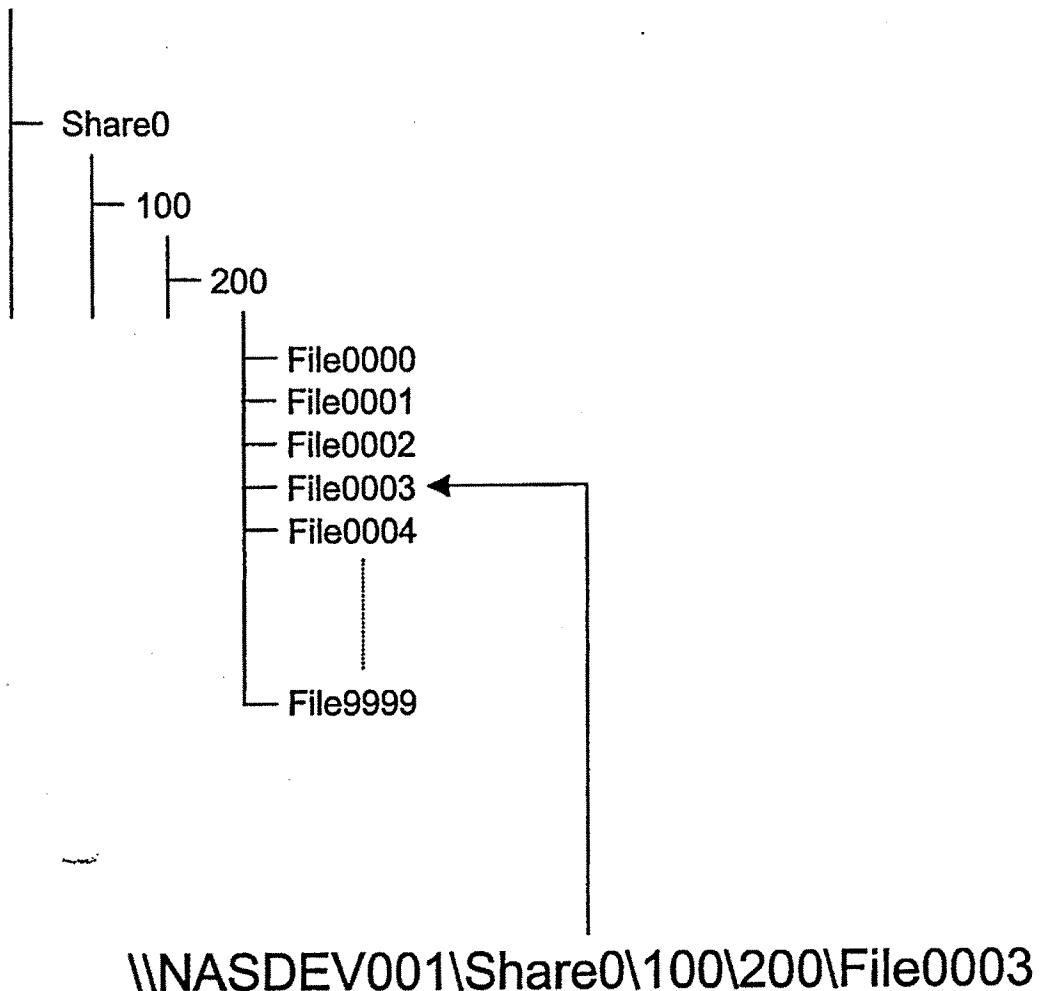
Figure 1 – Schematic Diagram of Components

Figure 2 – Directory Tree Structure of a Sample Storage Device

NASDEV001



3/6

Figure 3 – Sample Device Table

31 Device ID	32 Status	33 Content Type	34 Target Expiry Date	35 Actual Expiry Date	36 Kbytes Waiting	37 Kbytes Free	38 Current File
001	Full	1	31-Jul-04	03-Aug-04	0	0	4546
002	In Use	1	31-Jul-05	26-Jan-04	2048	51800	435
003	Off-Line	2	31-Jul-04	08-Aug-04	0	1800	7850
004	In Use	2	31-Jul-05	22-Jan-04	14	20900	283
005	Free	-	-	-	-	102400	0

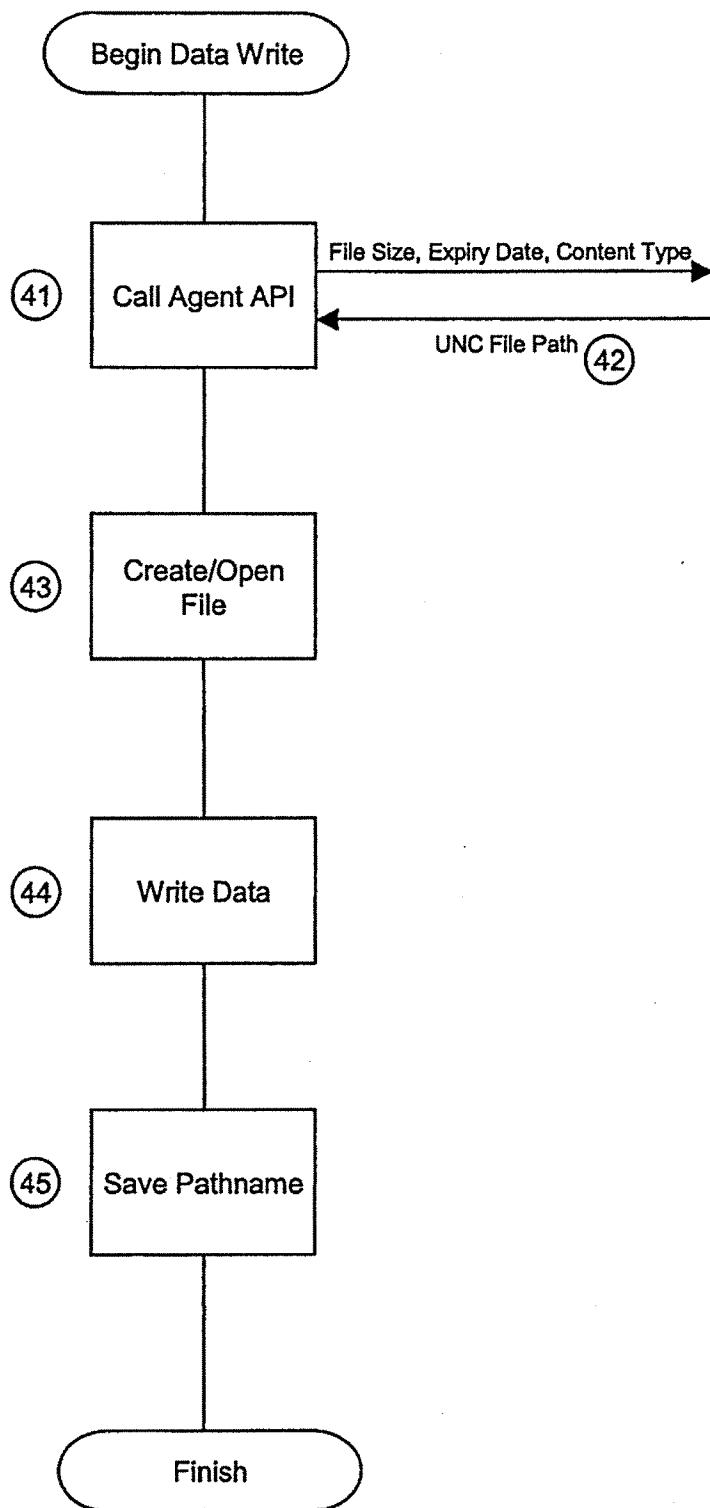
Figure 4 – Write Data Procedure

Figure 5 – Storage Device Selection

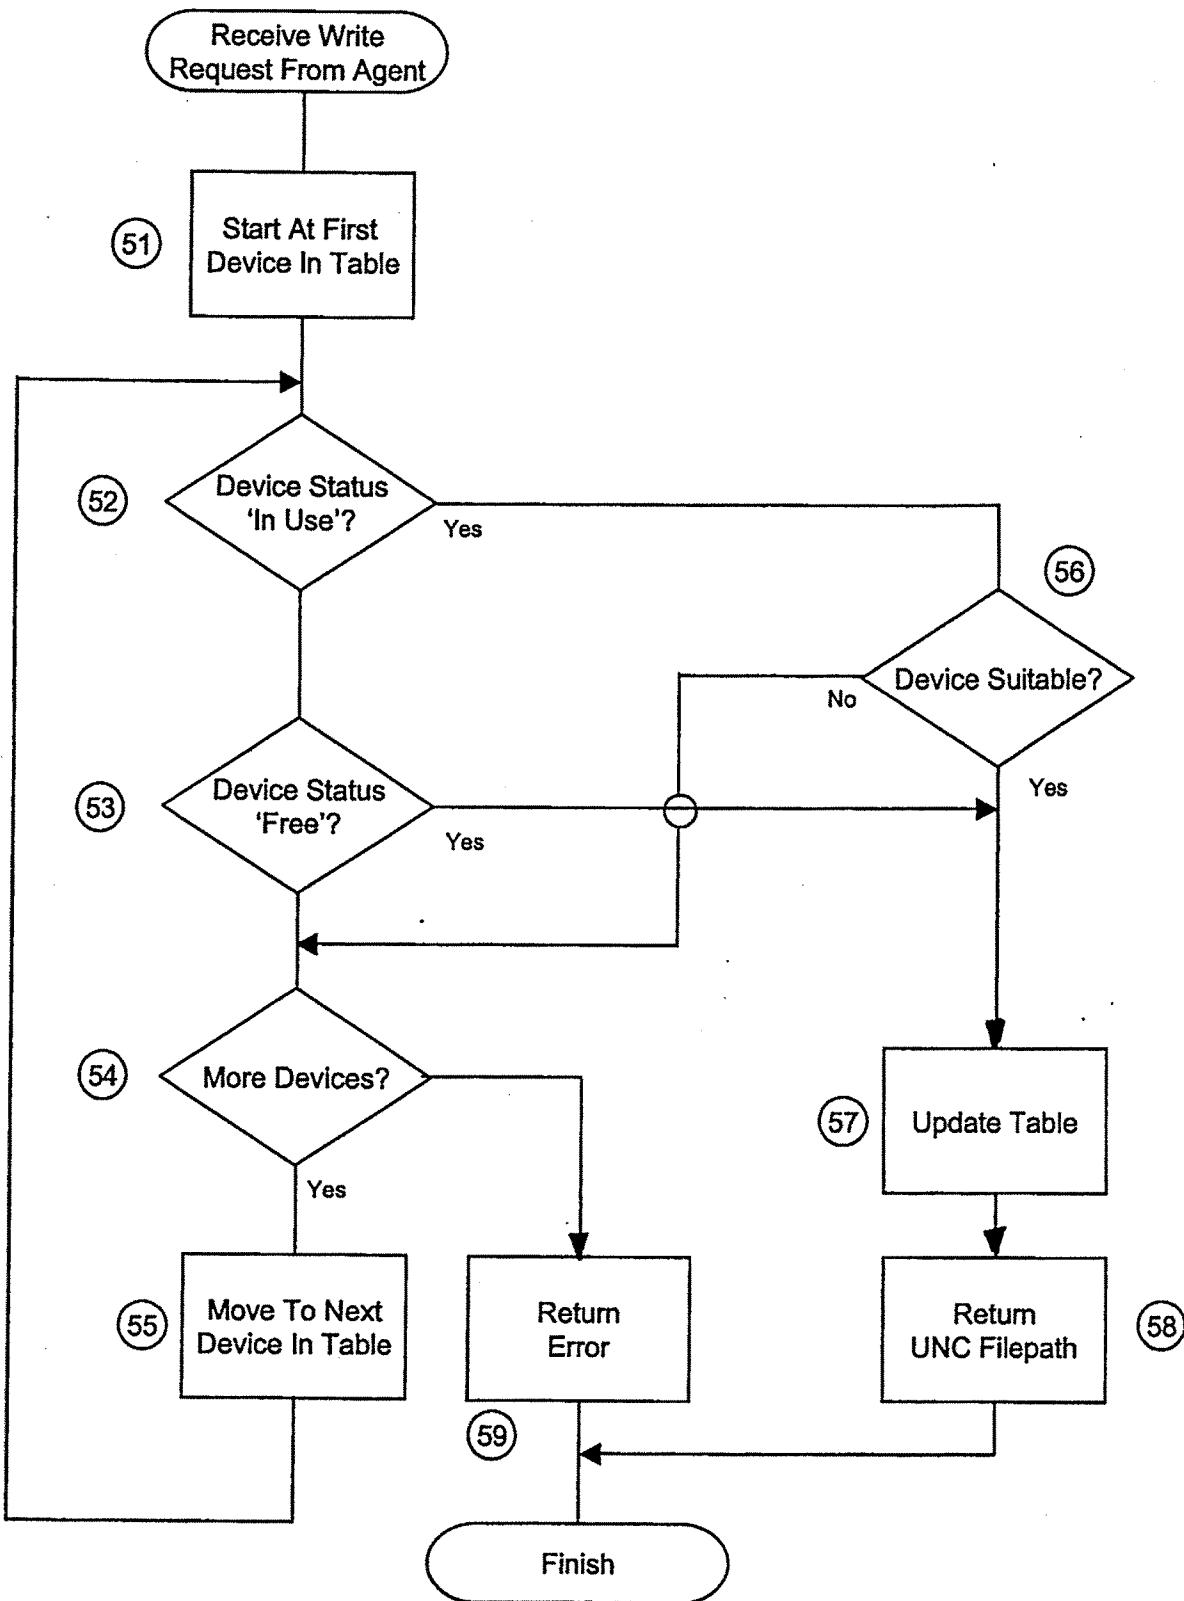
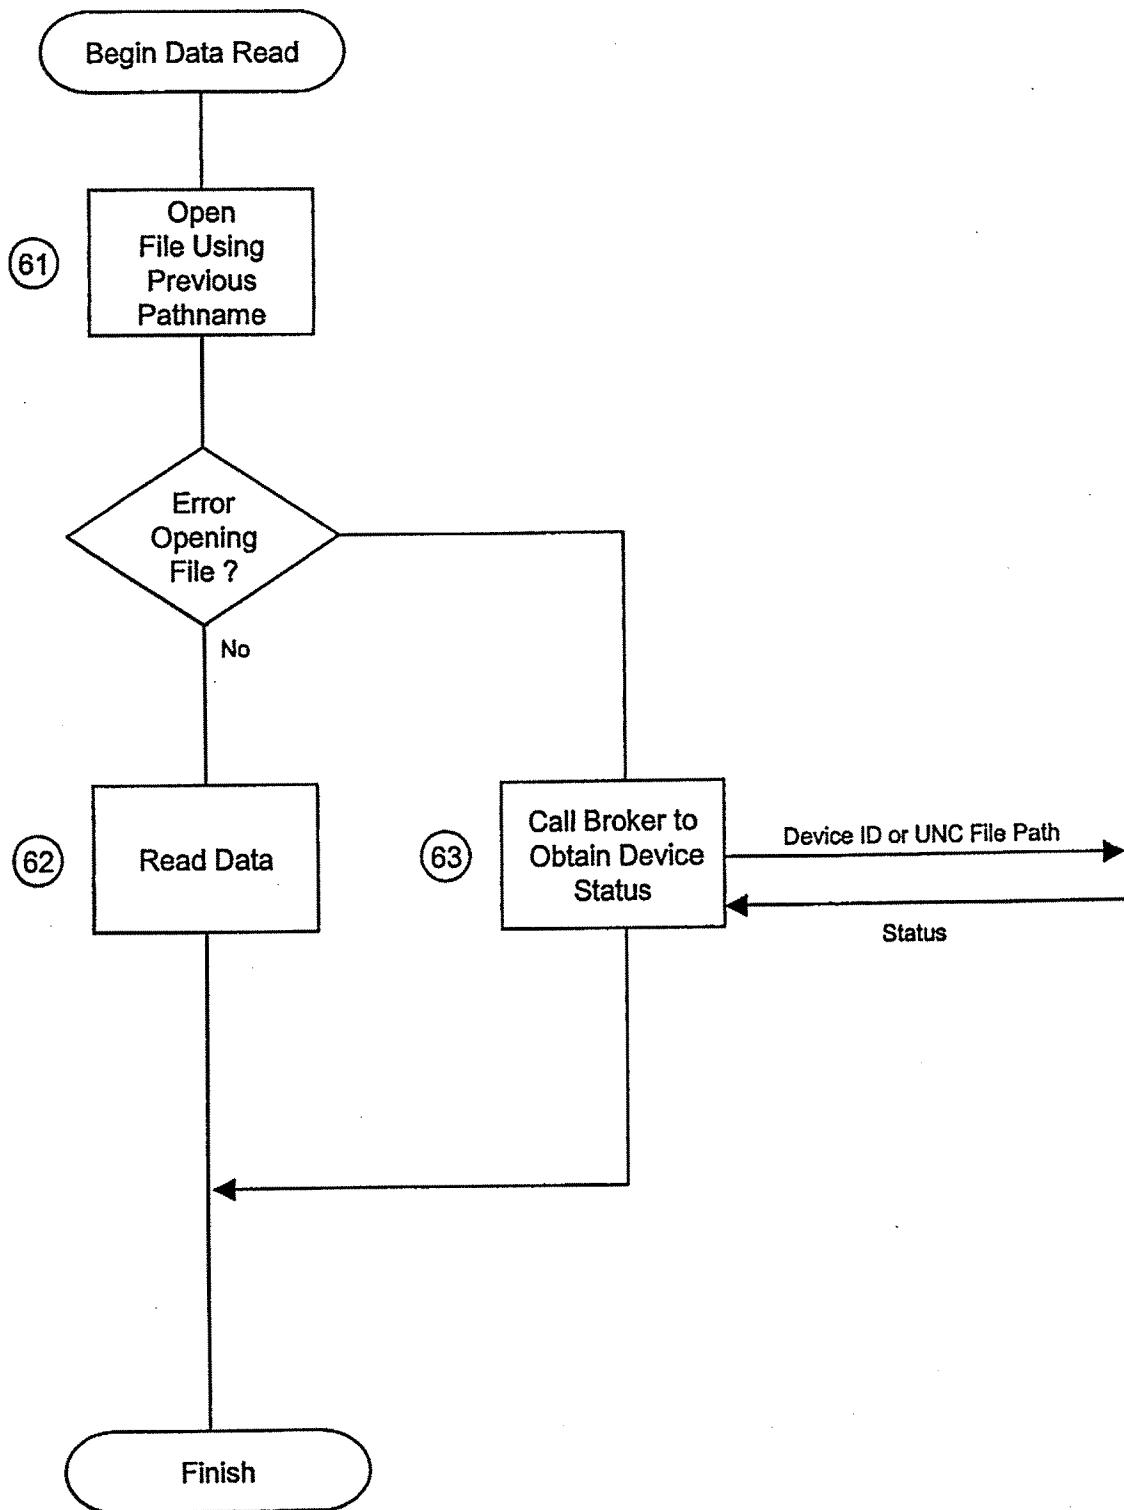


Figure 6 – Read Data Procedure

DATA STORAGE SYSTEM

Field of the Invention

This invention relates to a data storage system and data storage method, and more particularly to the long term archiving of electronic communications and other documents.

5

Background of the Invention

Recent increases in regulation, for example from bodies such as the Securities and Exchange Commission (SEC), National Association of Securities Dealers (NASD), their counterparts in other countries, and by legislation such as the Sarbanes-Oxley Act in the United States, have significantly increased the requirement to retain copies of electronic communications, such as electronic mail, instant messages, and other document types.

10

15

20

25

Coupled with the significant growth in the use of electronic communications and documents, the storage and management costs of retaining such data have escalated enormously, particularly for large organizations with many thousands of employees.

Recent studies show that a typical corporate employee currently handles around 7 megabytes of electronic mail per day. This number is forecast to increase to around 14 megabytes by 2007. Assuming an average 5 year retention requirement, this means that the storage needs of a 25,000 user organization could grow to over 300 terabytes, just for electronic mail.

In the past, much of the long term document retention needs of an organization have been satisfied using technologies such as magnetic tape and in some cases write-once optical devices. For many organizations however these are no longer viable options. This is because of an increasing need to be able to rapidly

30

retrieve individual documents and to be able to search the content of all stored documents in a matter of seconds. Slow access times mean that this is impossible in the case of magnetic tape, and impractical in the case of optical devices which also suffer from relatively low capacity.

5 Currently, the only viable option is high capacity on-line data storage, which is produced by a relatively small number of large scale storage manufacturers, costing in the region of US\$100,000 per terabyte. This is in 10 sharp contrast to the 'commodity' end of the market where disk drive prices have fallen sharply while capacities have increased. A 30 gigabyte capacity hard disk can be purchased for as little as US\$50, equivalent to around \$1,700 per terabyte. One must, however, be cautious when 15 making comparisons of this type, because large scale systems include infrastructure such as power supplies and controllers, and provide manageability and resilience, thus lowering the effective cost of ownership; however, such comparisons do illustrate the point that there is a 20 wide pricing gap between the two.

The industry generally has also been slow to recognize that the requirements for long term on-line archival storage are fundamentally different to those of 'normal' on-line storage applications. In particular, in 25 a traditional on-line storage environment, data 'read' activity generally exceeds 'write' activity by several, if not many, times. Consequently, on-line storage is usually optimized for reading, by the use of techniques such as caching.

30 In contrast, archiving, particularly of transient data such as email and instant messages, demands very high write performance in order to rapidly process the constant stream of incoming data - otherwise the messages have to either be delayed, or a temporary copy of them must be 35 made, placing additional demands on the messaging system.

Storage technology has evolved considerably in recent years, with three new storage technologies emerging.

5 Network Attached Storage (NAS) couples one or more hard disk drives with a processor, memory and network connection to provide an inexpensive server dedicated to data storage, which can be accessed by any other networked system.

10 NAS has become a commodity item, available from a wide range of manufacturers at a cost below US\$10,000 per terabyte. It has a relatively low capacity however, currently a maximum of 1 terabyte in a single physical device.

15 Storage Area Networking (SAN) adopts a different approach, providing a pool of physical storage devices connected by a dedicated high speed network. Application and file servers connect to this dedicated network and share the available storage pool, which can be easily expanded without needing to take servers off-line.

20 SAN is available from a number of manufacturers, but requires purpose designed high speed networking components, which increase its cost considerably over NAS. It supports a large pool of storage devices, but typically does not share logical storage volumes between servers.

25 Content Addressable Storage (CAS) is specifically designed for archiving static documents, ie those which do not change over their storage lifetime. Applications pass a copy of the document they want stored, and are returned a unique 'token' which the application uses to subsequently retrieve the document. CAS provides the advantage that multiple copies of the same document are automatically identified and only a single copy is stored.

30 Since CAS manages its own operating system its capacity is practically unlimited, but the requirement to calculate the unique token for every file written creates an overhead which limits write performance. In terms of capital cost however, CAS is currently at the top of the scale.

Summary of the Invention

The invention is defined by the independent claims to which reference should now be made. Advantageous features are set forth in the dependent claims.

5 The preferred system provides a storage scheme for applications requiring persistent data storage, that is storage of data that once written is kept secure for an extended period of time until it is no longer required. This is in contrast to ordinary data storage in which it
10 is assumed data read, write and erase operations are carried out more frequently.

15 The applications preferably operate on one or more computers connected by a computer network to one or more storage devices. The term computer is understood to mean a server, client or stand-alone machine.

20 The invention allows a practically infinite number of physically separate storage devices to be used as archival storage means by one or more application programs. Application data is organized among the devices such that files with a similar expiry date (that is the date beyond which the files are no longer required to be archived) are grouped together on the same physical device, such that after the expiry date the device used for such files may be erased and re-used. Optionally, files with particular
25 content types may be grouped on the same physical device, or among physical devices of the same type.

30 The invention allows the storage system to start small, with perhaps only a few physical storage devices, and to be expanded by adding additional devices over time. This avoids a large initial capital outlay. Since the capacity of each physical device can be different, the invention allows the latest technology devices to be used whenever additional storage devices are added. This allows the most advantageous price/capacity ratio to be
35 achieved.

5 The invention also provides means for load scalability, such that data may be read or written as multiple parallel streams between applications and devices providing practically unlimited performance, particularly in writing data to the storage means.

10 The result is a high capacity, high speed storage system able to meet the archiving needs of even the largest organizations, at a cost which is an order of magnitude less than current solutions, and which allows advances in storage technology to be taken full advantage of.

15 The invention also provides means for data to be secured, such that it cannot be deleted by an external process until its expiry date has been reached, an important requirement in regulatory environments.

Brief Description of the Drawings

The invention will now be described in more detail, by way of example, and with reference to the drawings in which:

20 **Figure 1** is a schematic diagram showing the composition of the preferred embodiment of the invention;

Figure 2 shows the directory tree structure of one of the storage devices and a sample of the corresponding identifier for a given stored file;

25 **Figure 3** shows a sample of a device table maintained by the Broker process;

Figure 4 is a flowchart showing the procedure for writing data to the storage means;

30 **Figure 5** is a flowchart showing how the selection of storage device is determined; and

Figure 6 is a flowchart showing the procedure for reading data from the storage means.

Detailed Description of the Preferred Embodiment

The preferred embodiment of the invention will now be described in more detail with reference to the drawings.

Figure 1 shows the components of the preferred embodiment and their connectivity. Application Programs 1 running on one or more computers 2, herein after called Application Servers, are connected via standard network components 3, preferably using switched gigabit Ethernet or a similarly high capacity technology, to a set of NAS devices 4 and to a Broker program or process 5 preferably running on a separate physical machine 6. Also provided are Agents 7 running on the Application Servers, which comprise executable code providing a standardized application programming interface (API) to allow the Application Programs to easily communicate with the Broker Application.

Each NAS device is accessible to each of the Application Programs using standard file sharing methods, such as those under the usual UNC naming convention, each NAS device having a unique identifier by which it is addressed. As shown in Figure 2, each NAS device has a conventional file system comprising a directory structure and files in the usual manner. Figure 2 illustrates the directory structure of a NAS device having the unique identifier 001 (shown as 'NASDEV001'). Files named 'File0000' to 'File9999' are stored in directory 200, which is itself a subdirectory of directory 100. Finally, directory 100 is contained within the top-level directory 'Share0'. By way of illustration the path name to one of the files (File 0003) is also shown: '\\\NASDEV001\Share0\100\200\File0003'.

Arranging files in directories and subdirectories is preferred, as it makes looking-up any given file more efficient. This is pertinent in the case of the preferred system, where reading and writing data speeds should be optimised.

The security provisions for ensuring controlled access to data may be provided by means of the usual username/password authentication, and/or by making the network 3 a private network accessible only to the 5 Application Servers and allowing the Application Programs to control access. Such security provisions are assumed to be known and so will not be discussed further.

The Broker Application maintains a Device table with each row of the table corresponding to a single NAS 10 device, and each column corresponding to various configuration and status conditions of the NAS devices. Figure 3 shows an example implementation of the Device table maintained by the Broker 5.

The first column 31 contains a Device ID such that 15 each storage device can be separately and uniquely referenced.

The next column 32 indicates the current status of each of the respective devices. The status may be 'Full', indicating that the storage capacity of the device has 20 been entirely used up; 'In Use' meaning that data has been written to the device but that storage space remains; 'Off-Line' meaning that the device has been temporarily or permanently disconnected and is not presently available to receive data; and 'Free' which indicates that the device 25 is empty and can be written to.

The Broker manages the device table such that all 'Free' devices appear after devices with another status. This includes devices which are being 'recycled' after 30 erasure of the data they contain and which is no longer needed. Thus, when all of the data in a 'Full' device is past its expiry date, the device status will be set to 'Free' and its entry in the Device table will be relocated to the bottom of the table. The data can be erased and the device made ready to receive new data.

35 The next column 33, in the table indicates a 'Content Type' for the data that is stored in the device. When an Application Program requests write permission from one of

the storage devices, it will indicate to the Broker process the content type of the data to be stored. The application Program therefore sets the content type for the data. The content types themselves may however be defined by a user or system administrator to reflect the types of data that are required to be stored. The content type '1' illustrated in Figure 3, may for example, indicate 'e-mail' data, and '2' may indicate word processor files, or html pages. Content type codes may also be used to subdivide data of a given type, such as word processor files for examples, by importance. The code '3' for example may indicate word processor files that are of particular importance, and should be saved for a longer period of time than regular word processor files.

Preferably, each device stores only one type of data content as this makes the recycling of the device more straightforward, and operation of the preferred system more efficient. However, it is also possible that different regions of the same device are used independently to store different respective content types.

The next column 34 in the Device table, entitled 'Target Expiry Date', records the date on which it is anticipated that erasure of the contents of the device, and reclamation of the storage space will take place. This date is preferably set by the user or system administrator of the storage system, as a means of exercising control over the storage of the data in the respective storage devices. Depending on the implementation chosen, the target expiry date may be a fixed, inflexible date, such that data which is required to be stored until after the target expiry date is not allowed to be stored in the respective device. Alternatively, it could be a flexible date in which case data that is required to be stored until after the target expiry date is allowed to be stored, providing the date until which it is required falls within a predetermined margin of the target date.

5 The 'Actual Expiry Date', shown in column 35, on the other hand, reflects the latest date until which an item of data is to be stored. If data is subsequently received that must be stored until a later date than the indicated 'Actual Expiry Date', the Actual Expiry Date is updated.

10 The next column 36, entitled 'Kilobytes Waiting' indicates how much data allocated to a particular storage device has yet to be physically written to the device. As there is a finite rate at which data can be written to each device (currently 5Mb per second for a typical NAS device), it is necessary to monitor the data feed to try and avoid any particular device being overloaded.

15 The column 37 entitled 'Kilobytes Free' records how much remaining device space there is to store data, and lastly the column 38 entitled 'Current File' indicates the current file number. As illustrated in Figure 2, this preferably ranges from 0 to 9999 and will be incremented by 1 each time a new file is written. Other naming conventions, or ranges of numbers are also possible.

20 There now follows a description of how data is stored and retrieved by the preferred system:

25 Figure 4 shows a flowchart of the data write process. When an Application Program 1 requires that a file be written to the storage means, it first makes a call, Step 41, to the Broker Application, via the Agent API 7. In making this call it preferably passes three parameters; the length of the file to be stored, the expiration date beyond which the file will no longer be required, namely 'The Actual Expiry Date' (or a null date indicating that the file is to be kept forever), and an indication of data or content type. This third parameter may be optional.

30 35 The Broker returns (Step 42) either a fully qualified pathname or a set of identifiers from which the fully qualified pathname can be constructed. According to the implementation, the file can be created (that is an entry made in the appropriate directory on the appropriate device) by the Broker and then subsequently opened by the

Application Program, or the file can be created and opened (Step 43) by the Application Program itself (as shown in Figure 4). The Application Program then immediately writes the data file to the said location (Step 44). The 5 Application Program then stores (Step 45) either the fully qualified pathname or the identifiers which are used to construct it, such that the file can be read back at a later stage.

Figure 5 shows a flowchart of the process the Broker 10 uses to determine which device a given file should be written to. On receipt of a request from an Agent (step 41 in Figure 4), the Broker consults (Step 51) its Device table (such as shown in Figure 3) starting with the first device in the table. The Broker first determines at Step 15 52 if the device has an 'In Use' status, that is it is available for the storage of new files. If it has this status then control passes to step 56 where the Broker determines if the current operating parameters are suitable to allow the file to be stored.

20 In making this determination, the Broker considers a number of factors. Firstly, that the intended content type of the data (passed as a parameter by the Agent) matches that of the device in the table. As mentioned earlier, the necessity for the data content type to match 25 that of the device can be waived depending on the implementation. Secondly, that the device has enough free space to hold the file (the file size being passed by the Agent as a parameter, and the table containing the current amount of free space on the device). Thirdly that the 30 target expiry date of the device is within a given time period of the required expiry date (passed as a parameter by the agent). Lastly that the amount of data waiting to be written to the device does not exceed a given threshold representing an overload condition of the device.

35 The content type is entirely determined by the Application Program. It is used to group content of a certain type together. This might, for example, be used

to keep mail from different departments on different devices. The content type field could also be used for other purposes, for example to differentiate different types of device, allowing more important files to be 5 stored on more resilient storage devices, for example keeping customer communications on devices with mirrored disk drives.

Various algorithms can be used in considering whether storage of data with a given required expiry date or a 10 given device is acceptable. At its most simple, the target date can be a hard expiry date set in advance by the system administrator, such that if the required expiry date occurs at any time after the target date then the device is considered unsuitable for the data. 15 Alternatively, an algorithm can determine that the data will be accepted by the device if the required expiry date is not more than, say, 3 months beyond the target expiry date.

Most advantageously, the algorithm considers multiple 20 factors, including the amount of free space on the device, and the rate at which it is filling with data (and hence the date it is likely to become full). The objective is to fill a given device with data files which have the 25 earliest possible expiry date, such that the device may be eligible for re-use as soon as possible, and without wasting space.

The overload threshold amount is typically the amount 30 of data that the device is able to process in a given time period. The frequency with which the 'Kilobytes waiting' field 36 is updated preferably matches this time period. A suitable allowance can be made for read activity; for example the threshold might be reduced to 80% of the device's throughput capacity.

If it is determined that the current operating 35 parameters of the device are suitable to allow the data to be stored, then control passes to step 57, where the device table is updated as follows. If the required

expiry date of the data being stored is later than the actual expiry date in the table, then the actual expiry date is replaced by it. The field containing the amount of data waiting to be written is increased by the size of
5 the file to be written (passed as a parameter by the Agent). The field containing the amount of free space on the disk is reduced by the size of the file to be written. For this purpose, the size of the file to be written may be rounded up to a multiple of the device's disk
10 allocation unit size in order to maintain a more accurate free space count. The current file number is then incremented.

The incremented file number is then used to construct a UNC file path, and this is returned to the Agent in step
15 58. It will be appreciated that while the table in Figure 3 shows a simple current file number, this could be used to construct a more complex path such as that described with reference to Figure 2. The purpose of a more complex path is to limit the number of files in a given directory,
20 since most operating systems become less efficient as the number of directory entries in a given directory increases. If new directories need to be created, then this can be done by the Broker at this stage.

If, in Step 56, the device is not found to be
25 suitable, such as if the Content Type does not match, insufficient storage space remains, or if allocating the data to the device will result in an overload condition, then control passes to Step 54, where a check for an alternative storage device is made. Assuming there is
30 another device, control passes to step 55 in which the next device in the device table is considered. The status of that device is then determined as before in Steps 52 to 54. A device pointer may be used in Steps 51 to 55 in conjunction with the device identifiers shown in column 31
35 to monitor the current device under consideration.

If no alternative device is found in step 54, then control passes to step 59 where an error is returned to

the Agent since no suitable device can be found on which to store the file.

5 If in Step 52, the device is not found to have an 'In-Use' status, then control passes to Step 53, where a check is made to see if the device is 'Free'.

10 If the device has a status of 'Free' then the table update in Step 57 is as follows. The status is changed from 'Free' to 'In-Use'. The content type is set to the requested content type (passed by the Agent as a parameter). The target expiry date is set either to the requested expiry date (passed by the Agent), or by some suitable algorithm to a date beyond it. The amount of data waiting is set to the size of the file to be written (again passed by the Agent), and the current file number 15 is set to 1. If the amount of free space on the drive has not already been set, then it is determined by the Broker and set in the table accordingly.

20 Control then passes to step 58 where the current file number is used to construct and return the UNC pathname as before.

25 As explained earlier, the Broker Application maintains the device Table such that devices with status 'Free' are added to the bottom of the table. This ensures that suitable devices which are already 'in-use' are used in preference to new, empty devices. This ensures that the available storage space is used efficiently.

30 If the device status is not determined as 'Free' in step 53, then control passes to step 54 where it is determined if there are any more devices in the table. If not, then control passes to step 59 where an error is returned to the Agent since no suitable device can be found on which to store the file.

35 In the case of a device being rejected because of an overload condition in the Kilobytes waiting column, the Broker may retry the overloaded device again after a predetermined delay to see if it has become available before issuing an error. An appropriate notification may

be given to the Application Program that requested write permission while this is occurring.

Otherwise, the device pointer is incremented (step 55) and control passes back to step 52, where the next 5 device in the table is considered.

It will be appreciated that new devices can be added to the system at any time by extending the Device table, adding an entry for the new device at the end of it, and setting the status for that device to 'Free'. The Broker 10 5 can be implemented to periodically check for new devices and update the device table automatically.

Another process that must be carried out by the Broker application is to periodically reduce the field containing the amount of data waiting to be written to 15 each device by the amount of data that can be written in such period. For example, if a given device is capable of accepting data at a rate of 5 megabytes per second (a typical performance rate) then the table can be updated once every second, reducing the amount in the data waiting 20 field by 5 megabytes. If this causes the data waiting amount to be negative, then it is set to zero as all data waiting will have been written. It will be appreciated that the update period and adjustment amount can be proportionately varied.

Because each hardware component in the preferred embodiment is connected via a network switch, multiple Application Programs each have a full capacity connection to each Storage device. Since the Broker avoids 25 30 overloading a device by moving write activity to another device, performance of the overall system can scale to much higher levels than with a single connection storage sub-system.

The Broker can also be configured to remove write and/or delete access to each file after a given time 35 interval. For example, 5 seconds after the Broker passes the UNC pathname back to the Agent it can change access control to the file on the device such that it is write

protected and cannot be deleted by any other process. This is achieved using a call to a standard operating system function on the device. While 5 seconds will likely be sufficient time for the Application to open the 5 file and begin writing data to it, in the case of a large file the Application Program may not have completed the task. However in the case of typical operating systems this does not cause a problem, since access rights are only checked at the time a file is opened. Alternatively, 10 the Broker can periodically check that the file has been created, written to, and closed before changing the access rights.

This method of locking the file ensures that it cannot be accidentally deleted by any process other than 15 by the Broker, which is configured not to do so until the expiry date for each file has been reached, when the Broker can either remove the deletion protection, or optionally delete the file itself. Thus the system effectively becomes a write once, read many (WORM) device, 20 at least until the expiry date of each file, and as such meets both the regulatory requirements for storage compliance (such as SEC and NASD rules) and the organization's need to free storage for re-use as soon as possible.

25 Figure 6 shows a flowchart of the data read process. When an Application Program needs to read data from a file it simply opens it Step 61 using the fully qualified pathname it stored at step 45 during the write process, constructing the pathname from the set of identifiers if 30 necessary. The Application Program then reads the data directly from the file, step 62. In the event the Application Program receives an error opening the file, for example if the file has been deleted after its expiry date, then it can consult the Broker Step 63 to determine 35 the status of the device and the cause of the error.

Although the preferred embodiment of the invention has been described as comprising a number of NAS devices,

it will be appreciated that any suitable storage devices, or combination of storage devices, could also be used.

Claims

1. A method of operating a data processing system, the system comprising one or more application programs requiring persistent data storage for data files of application data, a plurality of storage devices each accessible via a computer network to one or more computers executing the application programs, and a broker program, wherein the method comprises:

receiving, by means of the broker program, a request from an application program for storage of a data file of application data, the request including an expiry date, beyond which the data file is no longer required and may be deleted; and

selecting for the data file which of the storage devices will be used to store the data file in accordance with the characteristics of the application data to be stored, including the expiry date, and the state of the storage devices.

2. A method according to claim 1, comprising monitoring, by means of the broker program, the remaining storage space available on each of the storage devices, to distinguish between in-use storage devices which have had data files written to them and empty storage devices which have not.

3. A method according to claim 2, wherein the selecting step comprises selecting in-use storage devices in preference to empty storage devices.

4. A method according to claim 2, comprising:

monitoring how much data is waiting to be written to each storage device, to detect an overload condition in the process of writing the data; and

selecting, if an overload condition is detected for a storage device selected for storage, a different storage device for storage.

5. A method according to claim 4, wherein the selecting step comprises selecting in-use storage devices in preference to empty storage devices.

6. A method according to claim 1, comprising:
storing, for each storage device, the latest expiry date of data files stored on that device, or of data files that are to be stored; and
permitting application data to be stored on a storage device if its expiry date is within a predetermined range of the latest expiry date;

such that application data with similar expiry dates can be stored together and when such similar expiry dates have passed the storage device can be erased and re-used.

7. A method according to claim 6, comprising selecting another storage device for storage, if the expiry date of the application data is outside of the predetermined range of the latest expiry date.

8. A method according to claim 1, comprising storing for each of the storage devices a target expiry date, and selecting which of the storage devices to use in dependence on a comparison of the expiry date and the target expiry date.

9. A method according to claim 8, comprising preventing application data from being stored on a storage device, if the target expiry date for that storage device is earlier than the expiry date.

10. A method according to claim 8, comprising preventing application data from being stored on a storage device if the target expiry date for that storage device is earlier than the expiry date by more than a predetermined margin.

11. A method according to claim 1, comprising, after the data file has been written to the storage device, preventing the file from being modified or deleted until the expiry date has passed.

12. A method according to claim 6, comprising, after the latest expiry date has passed, erasing the contents of the storage device.

13. A method according to claim 1, wherein the characteristics of the data to be stored include a classification of the content of the application data.

14. A method according to claim 13, comprising:
storing for each of the storage devices, a target content type;
comparing the classification of the content of the application data and the target content type; and
preventing the application data from being stored in a storage device if the target content type for that device and the classification do not match.

15. A method according to claim 1, wherein the characteristics of the data to be stored include the application program which requires its storage.

16. A method according to claim 1, wherein the characteristics of the data to be stored include the size of the application data.

17. A method according to claim 1, comprising writing the application data to and/or reading the application data

from the storage devices directly by means of the application programs via the computer network.

18. A method according to claim 1, comprising notifying the application program of the storage device used to store the data file as determined by the broker program, such that the application program can store means to identify the device.

19. A method according to claim 18, wherein the data files are retrieved from the storage device by the application program directly via the computer network and without reference to the broker program.

20. A method according to claim 1, comprising determining, by means of the broker program, the directory location for storage of the data file on the storage devices.

21. A method according to claim 20, comprising creating, by means of the broker program, the directory entry for the data file in the directory location in anticipation of data being written to the file by the application program requesting storage.

22. A method according to claim 20, comprising notifying, by means of the broker program, the directory location of the data file to the application program for storage by the application program.

23. A method according to claim 20, wherein the directory entry for the data file in the directory location is created by the application program prior to data being written by it.

24. A method according to claim 1, wherein the state of the storage devices includes the current availability of such devices for data to be written thereto.

25. A method according to claim 1, wherein the state of the storage devices includes the amount of free space available in the storage devices.

26. A method according to claim 1, wherein the state of the storage devices includes the rate at which data is being read from and/or written to such devices.

27. A method according to claim 1, comprising monitoring the status of the storage devices, detecting when new storage devices have been added, and making these available for storage.

28. A method of operating a data processing system to store data, comprising:

receiving a request from an application program, among one or more application programs, for storage of a data file of application data, the request including an expiry date, beyond which the data file is no longer required and may be deleted;

determining one or more characteristics of the data file, including the expiry date;

monitoring the status of storage devices in a plurality of storage devices; and

selecting a storage device from the plurality of storage devices to store the data file, in accordance with the characteristics of the application data to be stored, including the expiry date, and the state of the storage devices;

wherein the data file is stored on the selected storage device.

29. A system for storing data comprising a broker program arranged to receive a request for storage of a data file of application data from one or more application programs requiring persistent data storage, the request including an expiry date, beyond which the data file is no longer required and may be deleted; and select for the data file which of a plurality of storage devices, accessible to the one or more application programs and the broker program via a computer network, will be used to store the data file;

wherein the Broker program selects the storage device in dependence on the characteristics of the data file to be stored, including the expiry date, and the state of the storage devices, such that the data file is stored in the allocated storage device.

30. A system according to claim 29, comprising a look-up table accessible by the broker program, wherein the look-up table stores the remaining storage space available on each of the storage devices such that in-use storage devices which have had data files written to them and empty storage devices which have not are distinguished from each other.

31. A system according to claim 30, wherein the broker program is operable to select in-use storage devices in preference to empty storage devices.

32. A system according to claim 30, wherein the look-up table stores how much data is waiting to be written to each storage device, such that an overload condition in the process of writing the data to a storage device can be detected; wherein the broker program is operable to select, if an overload condition is detected for a storage device selected for storage, a different storage device for storage.

33. A system according to claim 32, wherein the broker program is operable to select in-use storage devices in preference to empty storage devices.

34. A system according to claim 29, comprising a look-up table accessible by the broker program, wherein the look-up table stores, for each storage device, the latest expiry date of data files stored on that device, or of data files that are to be stored, and

wherein the broker program is operable to permit application data to be stored on a storage device if its expiry date is within a predetermined range of the latest expiry date, such that application data with similar expiry dates can be stored together and when such similar expiry dates have passed the storage device can be erased and re-used.

35. A system according to claim 34, wherein the broker program is operable to select another storage device for storage, if the expiry date of the application data is outside of the pre-determined range of the latest expiry date.

36. A system according to claim 29, comprising a look-up table accessible by the broker program, wherein the look-up table stores, for each storage device a target expiry date, and the broker program is operable to select which of the storage devices to use in dependence on a comparison of the expiry date and the target expiry date.

37. A system according to claim 36, wherein the broker program is operable to prevent application data from being stored on a storage device, if the target expiry date for that storage device is earlier than the expiry date.

38. A system according to claim 36, wherein the broker program is operable to prevent application data from being stored on a storage device if the target expiry date for that storage device is earlier than the expiry date by more than a predetermined margin.

39. A system according to claim 29, wherein the broker program is operable to prevent, after the data file has been written to the storage device, the file from being modified or deleted until the expiry date has passed.

40. A system according to claim 34, wherein the broker program is operable to erase, after the latest expiry date has passed, the contents of the storage device.

41. A system according to claim 29, wherein the characteristics of the data to be stored include a classification of the content of the application data.

42. A system according to claim 41, comprising a look-up table accessible by the broker program, wherein the look-up table stores a target content type; and wherein the broker program is operable to compare the classification of the content of the application data and the target content type, and prevent the application data from being stored in a storage device if the target content type for that device and the classification do not match.

43. A system according to claim 29, wherein the characteristics of the data to be stored include the application program which requires its storage.

44. A system according to claim 29, wherein the characteristics of the data to be stored include the size of the application data.

45. A system according to claim 29, wherein the broker program is operable to notify the application program of the storage device used to store the data file as determined by the broker program, such that the application program can store means to identify the device.

46. A system according to claim 29, wherein the broker program is operable to determine the directory location for storage of the data file on the storage devices.

47. A system according to claim 46, wherein the broker program is operable to notify the directory location of the data file to the application program for storage by the application program.

48. A system according to claim 29, wherein the state of the storage devices includes the current availability of such devices for data to be written thereto.

49. A system according to claim 29, wherein the state of the storage devices includes the amount of free space available in the storage devices.

50. A system according to claim 29, wherein the state of the storage devices includes the rate at which data is being read from and/or written to such devices.

51. A system according to claim 29, wherein the broker program is operable to monitor the status of the storage devices, detect when new storage devices have been added, and make these available for storage.

52. A system for storing data, comprising:
one or more application programs requiring persistent data storage for data files of application data;
a plurality of storage devices each accessible via a

computer network to one or more computers executing the application programs; and

a broker program arranged to receive a request from an application program for the storage of a data file of application data, the request including an expiry date, beyond which the data file is no longer required and may be deleted; and to select for the data file which of the storage devices will be used to store the data file in accordance with the characteristics of the data file, including the expiry date to be stored, and the state of the storage devices.

53. A system according to claim 52, comprising an Application Program Interface running on the one or more computers to pass commands to and from the broker program and the application program.

54. A computer program product for controlling a computer in a data storage system, the computer being operable to receive requests from one or more application programs, running on one or more computers, and requiring persistent data storage for data files of application data, and operable to monitor a plurality of storage devices,

the computer program product comprising a recording medium readable by the computer, having program code stored thereon which when executed on the computer configures the computer to perform the steps of:

receive a request for storage of a data file of application data from an application program, the request including an expiry date, beyond which the data file is no longer required and may be deleted; and

select for the data file which of the storage devices will be used to store the data file in accordance with the characteristics of the data file to be stored, including the expiry date, and the state of the storage devices.

55. A computer program product according to claim 54, which when executed on the computer configures the computer to monitor the remaining storage space available on each of the storage devices to distinguish between in-use storage devices which have had data files written to them, and empty storage devices which have not.

56. A computer program product according to claim 55, wherein the selecting step comprises selecting in-use storage devices in preference to empty storage devices.

57. A computer program product according to claim 55, which when executed on the computer configures the computer to monitor how much data is waiting to be written to each storage device, to detect an overload condition in the process of writing the data; and

select, if an overload condition is detected for a storage device selected for storage, a different storage device for storage.

58. A computer program product according to claim 57, wherein the selecting step comprises selecting in-use devices in preference to empty devices.

59. A computer program product according to claim 54, which when executed on the computer configures the computer to:

store, for each storage device, the latest expiry date of data files stored on that device, or of data files that are to be stored; and

permit application data to be stored on a storage device if its expiry date is within a predetermined range of the latest expiry date;

such that application data with similar expiry dates can be stored together and when such similar expiry dates have passed the storage device can be erased and re-used.

60. A computer program product according to claim 59, which when executed on the computer configures the computer to:

select another storage device for storage, if the expiry date of the application data is outside of the pre-determined range of the latest expiry date.

61. A computer program product according to claim 58, comprising storing for each of the storage devices a target expiry date, and selecting which of the storage devices to use in dependence on a comparison of the expiry date and the target expiry date.

62. A computer program product according to claim 61, which when executed on the computer configures the computer to prevent application data from being stored on a storage device, if the target expiry date for that storage device is earlier than the expiry date.

63. A computer program product according to claim 61, which when executed on the computer configures the computer to prevent application data from being stored on a storage device if the target expiry date for that storage device is earlier than the expiry date by more than a predetermined margin.

64. A computer program product according to claim 57, which when executed on the computer configures the computer to prevent, after the data file has been written to the storage device, the file from being modified or deleted until the expiry date has passed.

65. A computer program product according to claim 59, which when executed on the computer configures the computer to erase the contents of the storage device, after the latest expiry date has passed.

66. A computer program product according to claim 54, wherein the characteristics of the data to be stored include a classification of the content of the application data.

67. A computer program product according to claim 66, which when executed on the computer configures the computer to:

store for each of the storage devices, a target content type;

compare the classification of the content of the application data and the target content type; and

prevent the application data from being stored in a storage device if the target content type for that device and the classification do not match.

68. A computer program product according to claim 54, wherein the characteristics of the data to be stored include the application program which requires its storage.

69. A computer program product according to claim 54, wherein the characteristics of the data to be stored include the size of the application data.

70. A computer program product according to claim 54, which when executed on the computer configures the computer to notify the application program of the storage device used to store the data file as determined by the broker program, such that the application program can store means to identify the device.

71. A computer program product according to claim 54, which when executed on the computer configures the computer to determine the directory location for storage of the data file on the storage devices.

72. A computer program product according to claim 71, which when executed on the computer configures the computer to notify the directory location of the data file to the application program for storage by the application program.

73. A computer program product according to claim 54, wherein the state of the storage devices includes the current availability of such devices for data to be written thereto.

74. A computer program product according to claim 54, wherein the state of the storage devices includes the amount of free space available in the storage devices.

75. A computer program product according to claim 54, wherein the state of the storage devices includes the rate at which data is being read from and/or written to such devices.

76. A computer program product according to claim 54, which when executed on the computer configures the computer to monitor the status of the storage devices, detecting when new storage devices have been added, and making these available for storage.

77. A computer program product for controlling a computer in a data storage system, the computer program product comprising a recording medium readable by the computer, having program code stored thereon which when executed on the computer configures the computer to perform the steps of:

receiving a request from an application program, among one or more application programs, for storage of a data file of application data, the request including an expiry date, beyond which the data file is no longer required and may be deleted;

determining one or more characteristics of the application data, including the expiry date;

monitoring the status of storage devices in a plurality of storage devices;

selecting a storage device from the plurality of storage devices to store a data file of application data in accordance with the characteristics of the data to be stored, including the expiry date, and the state of the storage devices;

wherein the data file is stored on the selected storage device.

78. A method of storing data substantially as described herein and with reference to the drawings.

79. A system for storing data substantially as described herein and with reference to the drawings.

80. A computer program product for controlling a computer, in a system for storing data, substantially as described herein and with reference to the drawings.